

The transition from TeraGrid to XSEDE

OPERATIONS

The challenge

For a decade, the National Science Foundation's (NSF) TeraGrid project provided researchers and students across the country with an integrated network of high-performance computers, data resources and tools, and research facilities among 11 partner sites.

In 2011, NSF enhanced the foundation laid by the TeraGrid project through its successor, the Extreme Science and Engineering Discovery Environment (XSEDE). As with TeraGrid, the main focus for XSEDE is the user community. The XSEDE Operations team ensures that users have the hardware, tools, and technical support they need to conduct cutting-edge research across a multitude of domains.

What are the operational differences with XSEDE, and how will these changes increase researchers' effectiveness in their use of the supercomputing resources?

The goal

XSEDE Operations improves user productivity through enhanced ease of use; reliable hardware, tools and support; and enriched quality assurance. The Operations team will track metrics to gauge overall success and continually improve.

"XSEDE provides an opportunity for us to review operations—what works and what doesn't—and improve on the procedures," said Victor Hazlewood of the National Institute for Computational Sciences (NICS).

XSEDE Operations supports:

- Cybersecurity, including an XSEDE Certificate Authority, two-factor authentication, Campus Bridging with InCommon, and security auditing services.
- Data management and coordination, including coordinated data services such as XSEDE-wide parallel file systems, data movement and management, and a framework for distributed archival replication.
- Networking implementation and support, including interconnecting major distributed XSEDE components and facilitating end-to-end performance for users.
- Software maintenance and coordination, headed by the Software Testing and Deployment team (formerly Software Support) whose primary focus is to improve the reliability and deployment of new software components through extensive pre-production testing and installation readiness prior to handoff to Service Providers or Systems Operational Support for deployment.
- Automatic distributed accounting and account management.
- Systems operational support, including frontline user support, systems administration for all centralized XSEDE services and monitoring through the 24x7 XSEDE Operations Center.
- User services, including online information services, allocations, user engagement, and training.

USER SERVICES

The challenge

The National Science Foundation and U.S. citizens have made a sizeable investment in the nation's high-performance computing resources, infrastructure, and workforce. XSEDE staff makes sure these resources are used efficiently and effectively in support of science. To help users maximize their time on the computing systems, XSEDE provides training, consulting, advanced interfaces, and timely and accurate information through the XSEDE User Services organization.

This year, NSF undertook the challenge of enhancing the foundation laid by the TeraGrid project through its successor, XSEDE. As with TeraGrid, the main focus for XSEDE is the user community. The XSEDE User Services team ensures that users have the information, tools, and access to staff required to conduct cutting-edge investigations across a multitude of research domains.

How do we enhance the services available to users and create a more productive environment for them?

The goal

The goal of XSEDE User Services is to improve user productivity through training, user news, allocations management, the XSEDE user portal, and enriched tools. The User Services team will track metrics to gauge overall success and continually improve.

"Based on the lessons we've learned in the TeraGrid, we will work toward improving all aspects of a user's experience with our resources and services," said Chris Hempel, user services director, Texas Advanced Computing Center (TACC).

XSEDE User Services supports:

- Training, including face-to-face sessions at HPC centers and campuses across the country, and increasing numbers of webcasts, online training, and virtual workshops.
- Consulting, including assistance with accounts, access, and information on how to run on XSEDE's advanced computing resources.
- User interaction and interface, including a revamped portal with added functionality and unvetted account access.
- User guides, including a new standardized template that makes all guides consistent and easy to follow.
- Allocations process, determined via peer review at quarterly meetings of the NSF XSEDE Resource Allocation Committee (XRAC).
- Feedback, including a centralized location for comments: feedback@xsede.org.

For more about XSEDE User Services, see:
<https://www.xsede.org/user-services>

CAMPUS BRIDGING

The challenge

Researchers, educators, and students have access to a wealth of digital resources, from the desktop to the campus to large national programs like XSEDE.

Currently it can be challenging to move among these various “cyberinfrastructures.” A chemist may become an expert at using her campus-level resources, but feel as though she’s starting from scratch when she’s ready to scale up to use larger resources through XSEDE. And her challenges will not be precisely those faced by a colleague at another institution, because each campus infrastructure is different—and researchers are moving to XSEDE from many, many diverse local systems.

How do we lower adoption barriers and build bridges so computational scientists and engineers are empowered to use any of the resources at their disposal?

The goal

XSEDE’s Campus Bridging initiative makes it easier to leverage diverse computing resources and various cyberinfrastructures.

“The idea is to allow any researcher to use local cyberinfrastructure and XSEDE as though they were peripherals to their laptop,” said Craig Stewart of Indiana University.

Not just a project, a world view

Campus Bridging is more than a single project—it’s a way of thinking about access to digital resources. XSEDE Campus Bridging co-leads Craig Stewart of Indiana University and Jim Ferguson of the National

Institute of Computational Sciences will keep the entire XSEDE team thinking about this philosophy: Moving between local and national resources should not be an insurmountable challenge for researchers but instead should be a smooth path to increased productivity.

Campus Bridging efforts include:

- Working with the XSEDE Architecture team to pilot the use of a Global Federated File System, which enables communities to plug their file systems into an XSEDE-wide global file system that can be transparently accessed from XSEDE and non-XSEDE resources.
- Communicating XSEDE architecture information to campuses, and advising campus-level IT staff on configuring their local resources in ways that will ease the transition to XSEDE.
- Helping Campus Champions to effectively communicate with their campus colleagues.
- Serving as a conduit for feedback to XSEDE leadership.

For more about XSEDE Campus Bridging, see:
<https://www.xsede.org/campus-bridging>

CAMPUS CHAMPIONS

The challenge

A researcher compiles data on the desktop for days. A biologist struggles to analyze huge amounts of data. An engineering team runs a project on the local cluster but needs to find a more powerful solution.

These researchers, scientists, and engineers need help to move beyond the technology problems so they can solve the science problems. They are learning that the solution often comes in the form of Campus Champions, volunteers who are knowledgeable, committed staff people who serve as a source of information, guidance, and support for their institutions.

Campus Champions is a three-year-old program supported by the Extreme Science and Engineering Discovery Environment (XSEDE). XSEDE, the National Science Foundation-funded follow-on project to the NSF's TeraGrid, supports, extends, and expands the Campus Champions program to further push scientific discoveries via supercomputing resources.

In August 2011, the program added its 100th institution, and now more than 130 individual Champions are serving as conduits to supercomputing information and resources. Campus Champion Jeff Pummill of the University of Arkansas said the addition of the 100th institution is a testament to how making needed connections and providing benefits to the research community can cause an organization to grow quickly and build a solid reputation.

It can be difficult for researchers to know about all the resources available to them, especially at larger institutions. Campus Champion Kim Dillman of Purdue University sees herself as a "translator," putting her experience to work to bridge information gaps and help users with various levels of expertise make the most of the available hardware, software and user support.

The goal

Champions meet with researchers on a regular basis; network with and learn from other Champions via conferences and meetings; host meetings, orientations, seminars and events to increase awareness of supercomputing resources; provide assistance with allocations; and remain well-informed about resources — especially those available at the national level such as XSEDE.

Champions also demonstrate how XSEDE differs from TeraGrid. "XSEDE is more about the people and the use of the resources than it is about the hardware," Pummill said. "There is still an emphasis on hardware, and that's not going away," he added, "but XSEDE provides a greater emphasis on how to more effectively use the resources."

Jeff Gardner of the University of Washington is a Campus Champion and a researcher currently working on a project that simulates large chunks of the universe to study how galaxies form. "Having lived on the other side of the fence, one of the challenges was that things are always changing, and how do you keep the community aware of that?" he said. "At our institutions, most of us have multiple roles, and Campus Champions is one tool in our arsenal."

Whether a Champion is a systems administrator, IT faculty or staff, a computer science or math professor, the CIO, deputy CIO, or a researcher, each one volunteers because there is a need for the Champion's expertise. Champions know that every successful supercomputer project requires a certain level of knowledge, and having that information is especially critical when someone is just getting started. As Pummill said, "Currently there is no better way to make that connection than through the Campus Champions."

For more about XSEDE Campus Champions, see:
<https://www.xsede.org/campus-champions>

EXTENDED COLLABORATIVE SUPPORT SERVICES

The challenge

To bring the best available knowledge and skills to bear on computational science problems within the XSEDE user community, ECSS helps maximize science and engineering results and share the knowledge gained as widely as possible.

Through ECSS, expertise is available to support users in a wide range of areas, from performance analysis and petascale-optimization techniques to building Science Gateways and workflow systems. Users may request this in-depth support through the XSEDE allocations process. These requests are referred to ECSS staff, who follow up with users to establish a well-defined project plan for a collaboration lasting up to one year.

ECSS has evolved from the Advanced User Support function of the predecessor TeraGrid program, but it significantly extends those services, with greater emphasis on outreach and on the development, in particular, through a Novel and Innovative Projects focus, of projects in areas or from institutions and communities that can exploit advanced computing but haven't traditionally used it.

Leadership and goals

ECSS is co-led by Ralph Roskies, Pittsburgh Supercomputing Center, projects director; and Nancy Wilkins-Diehr, San Diego Supercomputer Center, communities director.

"ECSS staff work both with user groups in fields familiar with high-performance computing (HPC)," said Roskies "and with the XSEDE outreach team to reach user groups, communities, and digital services that are new to HPC."

ECSS communities projects benefit large groups of researchers, through the development of Science Gateways, the optimization of popular community codes, and a full training and outreach program. Wilkins-Diehr also leads a new ECSS initiative called Campus Champions Fellows, which she says, "will provide an extended user support experience for XSEDE Campus Champions, by pairing them with ECSS staff for extended work on a project, so that Champions can learn ECSS techniques and bring them back to their campuses."

For more about XSEDE

Extended Collaborative Support Services, see:
<https://www.xsede.org/ecss>

One constant in the TeraGrid-to-XSEDE transition is the significance of the education, outreach, and training (EOT) component. There will always be new projects and new scientists requiring assistance.

Through education, outreach, and training, including the efforts of Campus Bridging and Campus Champions, XSEDE intends to accomplish the following.

- Offer curriculum development workshops synchronously.
- Work with institutions to offer CS&E certificate and degree programs.
- Immerse students in training, internships, mentoring and recognition programs.

- Work with campus staff to bridge between campuses and XSEDE resources.
- Continue to expand the Campus Champions program to reach more campuses and users.

- Provide in-depth support to develop research projects among new users and underrepresented communities.
- Facilitate infrastructure coordination efforts among XSEDE and campuses through Campus Bridging.

- Offer more quality content to address user needs and expand scope.
- Deliver more training synchronously and asynchronously to expand scale.
- Offer a “Train the Trainers” program to expand scale.

<https://www.xsede.org/education-and-outreach>